Max-Margin based Discriminative Feature Learning

Changsheng Li, Qingshan Liu, *Senior Member, IEEE,* Weishan Dong, Xin Zhang, and Lin Yang

arXiv:1412.4863v1 [cs.LG] 16 Dec 2014

***Abstract*—In this paper, we propose a new max-margin based discriminative feature learning method. Specifically, we aim at learning a low-dimensional feature representation, so as to maximize the global margin of the data and make the samples from the same class as close as possible. In order to enhance the robustness to noise, a $l_{2,1}$ norm constraint is introduced to make the transformation matrix in group sparsity. In addition, for multi-class classification tasks, we further intend to learn and leverage the correlation relationships among multiple class tasks for assisting in learning discriminative features. The experimental results demonstrate the power of the proposed method against the related state-of-the-art methods.**

***Index Terms*—Feature learning, max-margin, correlation relationship, group sparsity**

I. Introduction

Data classification plays a key role in many practical applications [1], [2]. However, the real-world data, such as image data, often distributes in a high-dimensional space, which makes the computation cost expensive and might degrade prediction accuracies of classification models. To cope with this issue, a popular way is dimensionality reduction, which tries to project the data into a low-dimensional subspace with the least information loss [3]. From the technical point of view, dimensionality reduction can be divided into two categories: feature selection based methods and feature transformation based methods. The feature selection based methods aim at selecting the most informative feature subset from the original feature set directly. The feature transformation based methods intend to seek a transformation matrix to map the high-dimensional data into a low-dimensional space. According to the characteristics of transformation matrix, feature transformation techniques can be further divided into linear and nonlinear ones. In this paper, we focus on the linear techniques, due to its simplicity and effectiveness.

Many linear feature transformation methods have been proposed over the last decades. Principal Component Analysis (PCA) [4] and Linear Discriminant Analysis (LDA) [5] are two classical linear algorithms, which try to capture global Euclidean structure of the data. Since the practical data often lies on or close to an intrinsically low-dimensional manifold, volumes of approaches focus on how to preserve such structure in recent years. The representative approaches include isometric feature mapping (Isomap) [6], locally linear embedding (LLE) [7], Laplacian eigenmaps (LE) [8], locality preserving projection (LPP) [9], and neighborhood preserving embedding (NPE) [10]. Meanwhile, some work integrates global and local information during feature transformation [11]. However, the above methods do not consider that how to effectively connect with the classifier in the context of classification. In order to alleviate this limitation, the maximum margin projection (MMP) algorithm [12] took advantage of a binary support vector machine (SVM) [13] classifier to seek some hyperplanes that separated data points in different clusters with the maximum margin. The random projection algorithms [14], [15] tried to find some Gaussian random projection matrices to preserve the pairwise distances between data points in the projected subspace, which can be effectively combined with some classifiers, such as SVM. [16] proposed a framework to simultaneously learn a linear dimensionality reduction projection matrix and a margin-based classifier defined in the reduced-dimensional space. The main idea of the Maximum Margin Projection Pursuit (MMPP) algorithm intends to integrate optimal data embedding and SVM classification in a single framework in both the cases of bi-class and multi-class classification tasks [17].

In this paper, we propose a new Max-Margin based feature transformation method to Learn Discriminative Features for classification, called MMLDF. The proposed method tries to learn a low-dimensional feature space, where we not only try to make the global classification margin maximized, but also we expect the samples from the same class to be close as much as possible. In order to be robust to noise, we introduce a $l_{2,1}$ norm regularization term to make the transformation matrix in group sparsity, which plays the role of feature selection to some extent [18]. Additionally, in many real-world applications, there are often correlations among multiple classification tasks, and capturing such correlation relationships is helpful for learning discriminative features and designing classifiers [19]. In light of this point, we introduce another regularization term to capture the correlation relationships among multiple tasks. Extensive experiments are conducted on six publicly available datasets, and the experimental results demonstrate the effectiveness of the proposed method against the state-of-the-art methods.

The remainder of the paper is organized as follows. Section II gives the details of proposed method. The experimental results are reported and analyzed in Section III. And section IV concludes the paper.

II. Proposed Method

Let $\mathcal{X} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ denote a training data set, where $\mathbf{x}_i \in \mathbb{R}^d$ is the i-th data point and $y_i \in \{1, \ldots, \mathcal{K}\}$ represents the corresponding class label. Our goal is to seek a projection matrix $\mathbf{P} \in \mathbb{R}^{d \times r}$ that maps a d-dimensional input vector to an r-dimensional vector ($r < d$) by $\mathbf{z}_i = \mathbf{P}^T \mathbf{x}_i$. For the transformation matrix $\mathbf{P}$, $\mathbf{p}_i$ represents its i-th row, and P_{ij}

C. Li, W. Dong, X. Zhang, and L. Yang are with IBM Research-China, Beijing 100094, China. (email: lcsheng, dongweis, zxin, ylyang@cn.ibm.com).

Q. Liu is with the B-DAT Laboratory at the school of Information and Control, Nanjing University of Information Science and Technology, Nanjing 210014, China. (email: qsliu@nuist.edu.cn).

denotes the (i,j)-th entry of. $tr(\mathbf{P})$ denotes the trace of $\mathbf{P}$, and $\|\mathbf{P}\|_{2,1}$ denotes $l_{2,1}$ norm of $\mathbf{P}$ as:

$$\|\mathbf{P}\|_{2,1} = \sum_{i=1}^{d} \|\mathbf{p}_i\|_2 = \sum_{i=1}^{d} \sqrt{\sum_{j=1}^{r} P_{ij}^2} \quad (1)$$

A. Binary Classification: $\mathcal{K} = 2$

Margin maximization has been demonstrated to be a good principle applied to various learning methods [20], [21]. Among these methods, support vector machine (SVM) is a popular one that maximizes global margin of data points of different classes, but it only focuses on how to learn a max-margin based classifier. In this paper, we target at learning a low-dimensional feature space to further improve the performance of max-margin based classification, and we propose a unified framework to integrate feature learning and classification together. We not only intend to make the global classification margin maximized, but also we expect the samples of the same class to be close as much as possible in the subspace. Moreover, we also consider the issue that the data are often corrupted with noise and redundant information in practical scenarios. Therefore, we define the following objective function:

$$\begin{aligned}\min J(\mathbf{P},\mathbf{w},b) = &\frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i=1}^{n} l(\mathbf{w},b,\mathbf{P};\mathbf{x}_i,y_i) \\ &+\eta \sum_{i,j=1}^{n} \|\mathbf{P}^T\mathbf{x}_i - \mathbf{P}^T\mathbf{x}_j\|^2 A_{ij} + \lambda\|\mathbf{P}\|_{2,1} \quad (2)\end{aligned}$$

where C, η and λ are three trade-off parameters. $\mathbf{w}$ is the learned weight vector. $\mathbf{P}$ is the transformation matrix that we want to learn, and $\mathbf{P}^T\mathbf{x}_i$ is the new low-dimensional representation of $\mathbf{x}_i$. A_{ij} denotes the edge weight of the within-class graph. Here, either Heat kernel or Simple-minded can be used for weighting the edges. For similarity, we choose Simple-minded in our experiment, which is expressed as:

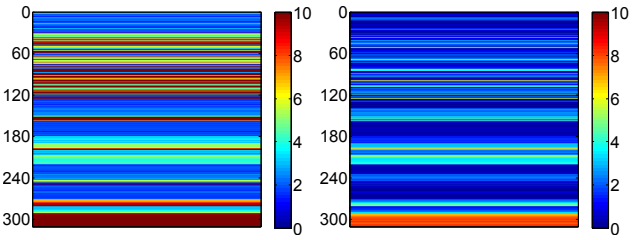
$$A_{ij} = \begin{cases} 1 & \text{if } y_i = y_j \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

$l(\bullet)$ is a hinge loss function. The standard hinge loss function in SVM is not differentiable everywhere. In order to take advantage of gradient-based optimization method to solve the proposed objective function, we adopt a similar but differentiable quadratic hinge loss:

$$l(\mathbf{w},b,\mathbf{P};\mathbf{x}_i,y_i) = [\min(0, y_i\cdot(\mathbf{w}^T\cdot(\mathbf{P}^T\mathbf{x}_i)+b)-1)]^2 \quad (4)$$

Minimizing the first two terms of (2) aims at finding a low-dimensional subspace, in which the margin can be maximized for classification. The third term seeks to make the scatter of the data in the same class as small as possible in the subspace. The last term is a regularization term to make the transformation matrix $\mathbf{P}$ sparse in rows, so that the learned subspace can be robust to noise, and the model complexity can be lowered too.

We take the LSVT Voice Rehabilitation dataset [22], a bi-class classification dataset, as an example to illustrate the effectiveness of the $l_{2,1}$ norm constraint on $\mathbf{P}$ in the objective function (2). Fig. 1 (a) and (b) represent the visualizations of $\mathbf{P}$ without $l_{2,1}$ norm and with $l_{2,1}$ norm in (2), respectively. We can see that many rows in $\mathbf{P}$ become sparse by adding the $l_{2,1}$ norm constraint, which can eliminate noisy or redundant features before feature transformation, and can reduce the model complexity.



(a) Without $l_{2,1}$ norm constraint (b) With $l_{2,1}$ norm constraint

Fig. 1. The visualization of the learned $\mathbf{P}$ ($r = 10$) on the LSVT Voice Rehabilitation dataset. Each row is the l_2 norm value of each row of $\mathbf{P}$. Dark blue denotes that the values are close to zero.

B. Multi-class Classification: $\mathcal{K} > 2$

In the case of multi-class classification, we can extend (2) to the following objective function:

$$\begin{aligned}&\min\ J_{mc}(\mathbf{W},\mathbf{b},\mathbf{P}) \\ &= \frac{1}{2}\sum_{m=1}^{\mathcal{K}} \|\mathbf{w}_m\|^2 + C\sum_{i=1}^{n}\sum_{m\neq y_i} l(\mathbf{w}_m,\mathbf{w}_{y_i},b_{y_i},b_m,\mathbf{P};\mathbf{x}_i,y_i) \\ &\quad + \eta\sum\nolimits_{i,j=1}^{n} \|\mathbf{P}^T\mathbf{x}_i - \mathbf{P}^T\mathbf{x}_j\|^2 A_{ij} + \lambda\|\mathbf{P}\|_{2,1} \quad (5)\end{aligned}$$

where $\mathbf{W} = [\mathbf{w}_1,\ldots,\mathbf{w}_{\mathcal{K}}]$ is the set of the learned weight vectors. $l(\mathbf{w}_m,\mathbf{w}_{y_i},b_{y_i},b_m,\mathbf{P};\mathbf{x}_i,y_i)$ measures the loss when the sample $\mathbf{x}_i$ is wrongly classified into the $m(m \neq y_i)$ class. Similarly, the loss function l is revised as:

$$\begin{aligned}&l(\mathbf{w}_m,\mathbf{w}_{y_i},b_{y_i},b_m,\mathbf{P};\mathbf{x}_i,y_i) \\ &= [\min(0,(\mathbf{w}_{y_i}^T\mathbf{P}^T\mathbf{x}_i + b_{y_i} - \mathbf{w}_m^T\mathbf{P}^T\mathbf{x}_i - b_m - 2)]^2 \quad (6)\end{aligned}$$

In real-world applications, one classification task is often correlated with other classification tasks, and mining the correlations among multiple tasks can be good for feature learning [19]. Our experiments also demonstrate this point (See the details in III-C). Thus, we add a regularization term into (5) to capture the correlation relationships among multiple classification tasks, and the new objective function becomes:

$$\begin{aligned}&\min\ J_{mc}(\mathbf{W},\mathbf{b},\mathbf{P},\Gamma) \\ &= \frac{1}{2}\sum_{m=1}^{\mathcal{K}} \|\mathbf{w}_m\|^2 + C\sum_{i=1}^{n}\sum_{m\neq y_i} l(\mathbf{w}_m,\mathbf{w}_{y_i},b_{y_i},b_m,\mathbf{P};\mathbf{x}_i,y_i) \\ &\quad + \eta\sum_{i,j=1}^{n} \|\mathbf{P}^T\mathbf{x}_i - \mathbf{P}^T\mathbf{x}_j\|^2 A_{ij} + \lambda\|\mathbf{P}\|_{2,1} + \rho\|\mathbf{W}^T\|_{\Gamma} \quad (7)\end{aligned}$$

where ρ is a trade-off parameter. $\|\mathbf{W}^T\|_{\Gamma} = tr(\mathbf{W}\Gamma\mathbf{W}^T)$ is the Mahalanobis norm of the matrix $\mathbf{W}^T$, where Γ plays the role of the inverse covariance matrix that encodes the correlation relationships between the weight vectors $\mathbf{w}_i$ [23]. The inverse matrix of Γ is constrained to be positive definite

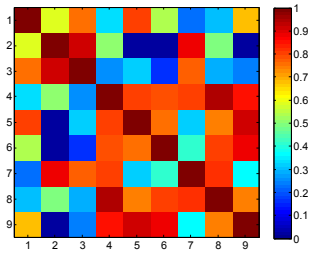


Fig. 2. The visualization of the learned correlation coefficient matrix on the Urban Land Cover dataset. Both x-axis and y-axis represent weight vectors. Dark red denotes that the correlation coefficients are close to 1.

and unit trace, in order to obtain a valid solution. The effect of the last term in (7) is to penalize the complexity of $\mathbf{W}$ relying on the Mahalanobis norm, as well as to learn the inverse covariance matrix Γ simultaneously.

We use the Urban Land Cover dataset [24], a multi-class classification dataset, to visualize the correlation coefficient matrix of weight vectors, which can be obtained based on the learned Γ. The result is shown in Fig. 2. We can see that there are indeed correlations among multiple classification tasks (e.g. the second and the third weight vectors).

C. Optimization Procedure

We first introduce how to solve the optimization problem in the binary classification case. The objective function (2) is not convex with respect to the variables $\mathbf{w}$, b, and $\mathbf{P}$ simultaneously. Therefore it is unrealistic to expect an algorithm to easily find the global minimum of J. Therefore, we adopt an alternating optimization strategy that is popular in machine learning field [25]. Under this scheme, we update $\mathbf{P}$, $\mathbf{w}$, and b in an alternating manner.

Update $\mathbf{P}$, with fixed $\mathbf{w}$ and b: When $\mathbf{w}$ and b are fixed, (2) becomes a convex problem, so $\mathbf{P}$ can be obtained by minimizing the following objective function:

$$\begin{aligned} J_1(\mathbf{P}) =& C\sum_{i=1}^{n}[\min(0, y_i \cdot (\mathbf{w}^T \cdot (\mathbf{P}^T\mathbf{x}_i)+b)-1)]^2 \\ &+\eta \sum_{i,j=1}^{n} \|\mathbf{P}^T\mathbf{x}_i-\mathbf{P}^T\mathbf{x}_j\|^2 A_{ij} + \lambda\|\mathbf{P}\|_{2,1} \end{aligned}$$

Taking a derivative of J_1 with respect to $\mathbf{P}$, we can obtain:

$$\begin{aligned} \frac{\partial J_1(\mathbf{P})}{\partial \mathbf{P}} =& 2C \sum_{(\mathbf{x}_i,y_i)\in\Theta} (\mathbf{x}_i\mathbf{x}_i^T\mathbf{P}\mathbf{w}\mathbf{w}^T - (y_i-b)\mathbf{x}_i\mathbf{w}^T) \\ &+ 2\eta\mathbf{X}\mathbf{L}\mathbf{X}^T\mathbf{P} + 2\lambda\mathbf{D}_P\mathbf{P} \end{aligned} \tag{8}$$

where $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_n]$. $\mathbf{L}$ is the Laplacian matrix. $\mathbf{D}_P$ is a diagonal matrix with the i-th diagonal element $\mathbf{D}_P(i,i) = \frac{1}{2\|\mathbf{p_i}\|_2}$. Θ denotes the set of $(\mathbf{x}_i, y_i)$ satisfying the following condition: $y_i(\mathbf{w}^T(\mathbf{P}^T\mathbf{x}_i) + b) - 1 \leq 0$.

Setting the derivative in (8) to zero, these is no closed-from solution of $\mathbf{P}$. Therefore, we adopt a gradient-based method to derive the optimal $\mathbf{P}$. Here we choose the limited-memory BFGS (L-BFGS) algorithm for its efficiency [26], [27], which is summarized in Algorithm 1.

Algorithm 1 L-BFGS

Input: Starting point $\mathbf{x}_0$, an integer $m > 0$, and a symmetric and positive definite matrix $\mathbf{H}_0$
$k \leftarrow 0$
Output: $\mathbf{x}_k$
Repeat
Computing $d_k \leftarrow -\mathbf{H}_k\nabla f_k$ using a two-loop recursion;
Computing $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k + \alpha_k d_k$, where α_k satisfies the Wolfe conditions;
$\mathbf{s}_k \leftarrow \mathbf{x}_{k+1} - \mathbf{x}_k, \mathbf{y}_k \leftarrow \nabla f_{k+1} - \nabla f_k$;
$\widehat{m} \leftarrow \min\{k, m-1\}$;
Updating Hessian matrix $\mathbf{H}_k$ using the pairs $\{\mathbf{y}_j, \mathbf{s}_j\}_{j=k-\widehat{m}}^{k}$;
$k \leftarrow k + 1$;
Until Convergence criterion satisfied.

Update $\mathbf{w}$, with fixed b and $\mathbf{P}$: When b and $\mathbf{P}$ are fixed, (2) is then convex in terms of $\mathbf{w}$, so we minimize the following objective function to obtain the optimal $\mathbf{w}$:

$$J_2(\mathbf{w}) = \frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i=1}^{n}[\min(0, y_i \cdot (\mathbf{w}^T \cdot (\mathbf{P}^T\mathbf{x}_i)+b)-1)]^2$$

Taking a derivative of $J_2(\mathbf{w})$ with respect to $\mathbf{w}$, and setting it to zero, we obtain:

$$\mathbf{w} = \begin{cases} (\mathbf{I} + 2C \sum\limits_{(\mathbf{x}_i,y_i)\in\Theta} \mathbf{P}^T\mathbf{x}_i\mathbf{x}_i^T\mathbf{P}\mathbf{w})^{-1} \\ \quad \times (2C \sum\limits_{(\mathbf{x}_i,y_i)\in\Theta} (y_i - b)\mathbf{P}^T\mathbf{x}_i) & \text{if } (\mathbf{x}_i, y_i) \in \Theta \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

Update b, with fixed $\mathbf{w}$ and $\mathbf{P}$: When $\mathbf{w}$ and $\mathbf{P}$ are fixed, (2) is then convex in terms of b, so b can be acquired by optimizing the following objective:

$$\min J_3(b) = \sum\nolimits_{i=1}^{n} l(b; \mathbf{w}, \mathbf{P}^T\mathbf{x}_i, y_i)$$

Taking a derivative of $J(b)$ with b, and setting it to zero, we obtain:

$$b = \begin{cases} \frac{\sum_{(\mathbf{x}_i,y_i)\in\Theta} (y_i - \mathbf{w}^T\mathbf{P}^T\mathbf{x}_i)}{|\Theta|} & \text{if } \Theta \neq \phi \\ b & \text{otherwise} \end{cases} \tag{10}$$

where $|\Theta|$ denotes the size of the set Θ.

The procedure of the proposed algorithm can be summarized in Algorithm 2.

Similarly, in the scenarios of multi-class classification, the procedure of updating the variables $\mathbf{P}$, $\mathbf{w}_i$, and b in (7) are the same as in the binary classification case. The rule for updating the variable Γ in (7) is as follows:

$$\Gamma = \frac{(\mathbf{W}^T\mathbf{W})^{1/2}}{tr((\mathbf{W}^T\mathbf{W})^{1/2})} \tag{11}$$

Details of the proof on (11) can be found in [23].

D. Time Complexity Analysis

The time complexity of Algorithm 2 consists of three parts: initialization on line 1, Laplacian matrix construction on line 2, and the iterative update of three variables on lines 3-8. The complexities of the first two parts can be ignored compared

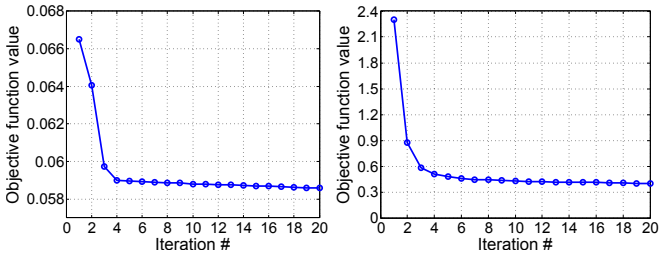
Algorithm 2 Max-Margin based Discriminative Feature Learning

Input: Training dataset $\mathcal{X} = \{\mathbf{x}_i, y_i\}_{i=1}^n$;
The parameters: C, λ, η, ρ;
Reduced dimension r

Method

1. Initialize iteration step $t = 0$; Randomly initialize $\mathbf{w}^t, b^t, \mathbf{P}^t$;
2. Construct Laplacian matrix $\mathbf{L}$;
3. **Repeat**
4. Fixing $\mathbf{P}^t$ and b^t, update $\mathbf{w}^{t+1}$ by Eq. (9);
5. Fixing $\mathbf{P}^t$ and $\mathbf{w}^{t+1}$, update b^{t+1} by Eq. (10);
6. Fixing $\mathbf{w}^{t+1}$ and b^{t+1}, update $\mathbf{P}^{t+1}$ by Algorithm 1;
7. $t = t+1$;
8. **Until** Convergence criterion satisfied.

Output: Transformation matrix $\mathbf{P} \in \mathbb{R}^{d\times r}$



(a) LSVT Voice Rehabilitation (b) Urban Land Cover

Fig. 3. Empirical study on the convergence of MMLDF on the LSVT Voice Rehabilitation dataset and the Urban Land Cover dataset.

to the third part. In the third part, we need to update $\mathbf{w}$, b, and $\mathbf{P}$, respectively. For updating $\mathbf{w}$, the worst complexity is $O(ndr + r^3)$. Updating b costs $O(ndr)$. For updating $\mathbf{P}$, it needs $O(t_1 * mdc)$ for the two-loop recursion scheme, where t_1 denotes the total iteration times. According to (8), the worst case of computing partial gradient w.r.t. $\mathbf{P}$ is $O(t_2 * (nd^2 + n^2d))$, where t_2 is the total number of computing the gradient. The complexity of evaluating the objective function values is $O(t_3*(rn^2+ndr+r^2d))$, where t_3 denotes the total evaluation times. Therefore, The total time complexity of MMLDF is of order $O(t * (ndr + r^3) + t_1 * mdc + t_2 * (nd^2 + n^2d) + t_3 * (rn^2 + ndr + r^2d))$, where t is the total iteration times in Algorithm 2. Since $t < t_3$ and $r < d$, the complexities of the parts updating $\mathbf{w}$ and b can be ignored, compared to that of updating $\mathbf{P}$, i.e., the time complexity of MMLDF is dominated by updating $\mathbf{P}$.

E. Evaluation of Convergence Rate

Although the convergence of the MMLDF algorithm cannot be proved theoretically, we find that it converges asymptotically in our experiments. Fig. 3 shows the convergence curves of MMLDF on the LSVT Voice Rehabilitation dataset and the Urban Land Cover dataset, respectively. As in Fig. 3, we can see MMLDF has a good convergence. It will converge after only about 6 iterations.

III. Experiments

A. Datasets and Experimental Settings

We evaluate the performance of MMLDF on six real-world datasets, including one aerial image dataset Urban Land Cover

TABLE I
Summary of experimental datasets. #Size, #Train, #Test, #Dim, and #Cat denote the number of samples, the number of training, the number of testing, the number of features, and the number of categories, respectively.

Dataset	#Size	#Train	#Test	#Dim	# Cat
Urban Land Cover	168	42	126	148	9
CNAE-9	1080	100	980	857	9
DNA	2,000	100	1900	180	3
Glioma	50	10	40	4434	4
LSVT Voice Rehab.	126	10	116	309	2
Epsilon	5,000	1000	4000	2000	2

[24], two biomedical area datasets DNA and Glioma, one large scale learning competition dataset Epsilon, one speech signal processing area dataset LSVT Voice Rehabilitation [22], and one business area dataset CNAE-9. The datasets DNA and Epsilon are downloaded from LIBSVM official web page[1], and the dataset CNAE-9 is downloaded from UCI Machine Learning Repository[2]. Datasets from different areas serve as a good test bed for a comprehensive evaluation. Table 1 summarizes the details of the datasets used in the experiments.

To verify the effectiveness of MMLDF, we compare it with the following seven related linear feature transformation methods:

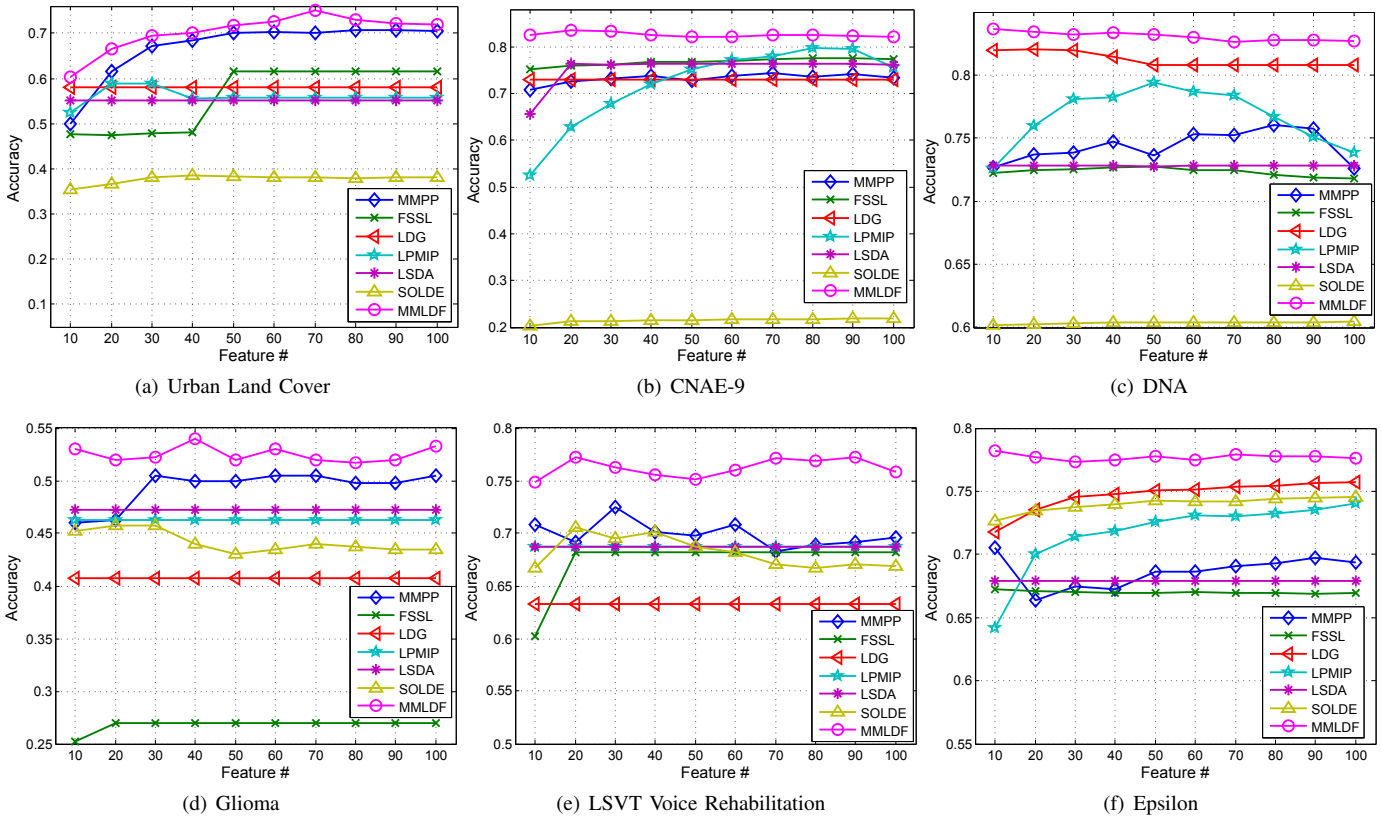
- SOLDE: Stable Orthogonal Local Discriminant Embedding [28] reduces the dimensions by considering both the diversity and similarity.
- LDG: Local Discriminant Gaussian [29] exploits a smooth approximation of the leave-one-out cross validation error of a quadratic discriminant analysis classifier[3].
- LPMIP: Locality-Preserved Maximum Information Projection [11] aims to preserve the local structure while maximizing the global information simultaneously.
- LapLDA: Laplacian Linear Discriminant Analysis [30] presents a least squares formulation for LDA, which intends to preserve both of the global and local structures.
- LSDA: Locality Sensitive Discriminant Analysis [31] tries to seek a projection which maximizes the margin between data points from different classes at local areas.
- FSSL: This method proposes a framework for joint feature selection and subspace learning [32].
- MMPP: Maximum Margin Projection Pursuit [17] aims to find a subspace based on maximum margin principle.

In the experiments, we vary the reduced dimensions from 10 to 100 with a stepsize of 10. After obtaining the low-dimensional feature representations using the above feature transformation methods, we use SVM with the linear kernel as the base classifier to evaluate the performance of the model. The classification accuracy is used as the evaluation measure. In our work, there are four parameters: C and λ, η, and ρ. The parameters C, η, and ρ are chosen by cross-validation, and the parameter λ is always set to 10^{-4} (We found when $\lambda = 10^{-4}$, the performance was consistently good on all the datasets). For

[1]http://www.csie.ntu.edu.tw/ cjlin/libsvmtools/datasets/
[2]http://archive.ics.uci.edu/ml/datasets/CNAE-9
[3]The MATLAB code for LDG was obtained from the authors of [29]

TABLE II
CLASSIFICATION ACCURACY(MEAN+DEVIATION (%)) OF DIFFERENT ALGORITHMS ON ALL THE SIX DATASETS

Datasets	Method							
	MMPP	FSSL	LDG	LPMIP	LSDA	SOLDE	LapLDA	MMLDF
Urban Land Cover	70.6±3.9	61.6±6.6	57.9±5.5	58.8±5.8	55.2±6.8	38.5±3.8	57.9±4.8	**75.2**±4.0
CNAE-9	74.4±2.9	77.5±3.0	72.9±2.8	79.8±2.7	76.4±6.9	21.8±3.7	78.2±4.1	**83.6**±1.6
DNA	76.0±2.5	72.7±1.6	82.0±2.8	79.4±2.9	72.8±1.9	60.5±3.5	74.9±1.8	**83.6**±1.6
Glioma	50.5±6.3	27.0±12.8	40.8±8.2	46.3±7.9	47.3±6.0	45.8±5.3	49.3±6.1	**54.0**±7.3
LSVT Voice Rehab.	72.5±6.1	68.2±5.5	63.3±8.3	68.7±6.3	68.7±8.0	70.6±8.8	72.4±7.6	**77.2**±7.2
Epsilon	70.6±1.3	67.3±0.8	75.7±0.8	74.1±1.0	67.9±0.8	74.6±0.4	70.4±0.3	**78.2**±0.6



(a) Urban Land Cover (b) CNAE-9 (c) DNA

(d) Glioma (e) LSVT Voice Rehabilitation (f) Epsilon

Fig. 4. Classification accuracy of different algorithms vs. the reduced dimensions on all the six datasets.

a fair comparison, the parameters in FSSL, LPMIP, LapLDA, LSDA, and MMPP are searched in the same space with that of MMLDF. For all the experiments, we repeat them 10 times, and report the average results.

B. General Performance

We first evaluate the classification performance of the proposed approach on all the data sets. Table II lists the best experimental results of all the algorithms. It can be seen that MMLDF consistently outperforms the other seven algorithms on all the six datasets. Compared with the second best result on each dataset, our method achieves 6.5%, 4.8%, 2.0%, 6.9%, 6.5%, and 3.3% relative improvement on the Urban Land Cover, CNAE-9, DNA, Glioma, LSVT Voice Rehabilitation, and Epsilon datasets, respectively. Some baselines obtain considerably poor performance on certain datasets (e.g., FSSL on Glioma, SOLDE on Urban Land Cover and CNAE-9). The reason may be that the generalization ability of these algorithms is limited, making them difficult to apply to different area datasets.

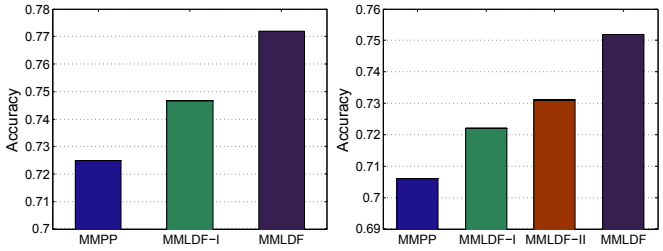
We also study the influence on the performance of different dimensions. Since LapLDA can be only reduced to $\mathcal{K}$ dimensions, where $\mathcal{K}$ denotes the number of the class, we do not compare our method with LapLDA. Fig. 4 shows the results. It can be seen that our method outperforms the other algorithms under all the cases. In addition, based on the results on the six datasets, we can see MMLDF is not sensitive to feature dimensions in a wide range.

C. Analyses on Components' Roles

We verify the effectiveness of the components in the objective functions (2) and (7), respectively. When the parameters λ, η, and ρ are set to zeros, MMLDF is reduced to MMPP, thus we use MMPP as the baseline. We perform the experiments on the binary class dataset LSVT Voice Rehabilitation and the multi-class dataset Urban Land Cover. The experimental

setting is as follows: we first set η to zero in (2), and set η and ρ to zeros in (7), in order to demonstrate the effectiveness of the module filtering noise and redundant features. We name it MMLDF-I for short. Then, we set ρ to zero in (7), which indicates that we learn the feature representation without considering multi-class correlations in the case of multi-class classification. We name it MMLDF-II.

The results are shown in Fig. 5. We can see that MMLDF-I outperforms MMPP on the datasets, which shows that the row sparseness constraint on the projection matrix is benefit for learning discriminative feature representation. On the LSVT Voice Rehabilitation dataset, MMLDF achieves better result than MMLDF-I, which means that with-class scatter minimum is good for classification. This point is also verified on the Urban Land Cover dataset, since MMLDF-II is superior to MMLDF-I. On the multi-class dataset, the result of MMLDF is better than that of MMLDF-II, which illustrates that capturing the correlations among multi-class classification tasks is helpful for enhancing the discriminative ability of the learned features. MMLDF achieves the best results on both of the two datasets. It shows that the combination of these components is effective for classification.



(a) LSVT Voice Rehabilitation (b) Urban Land Cover

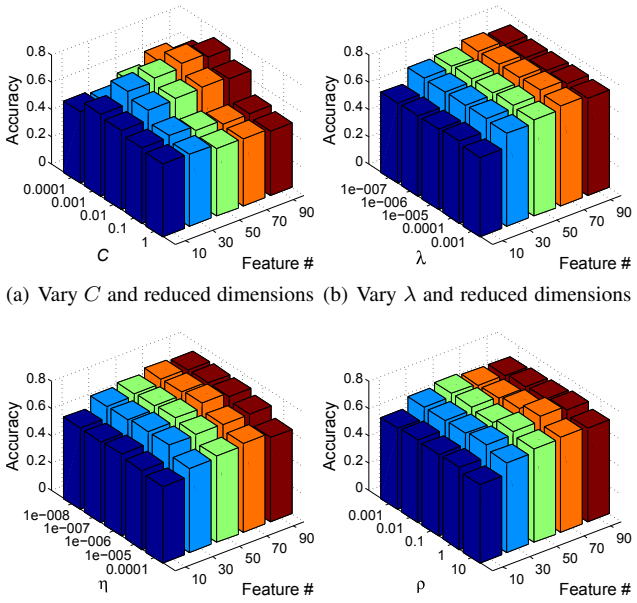
Fig. 5. Verify the effectiveness of each component in our algorithm on the LSVT Voice Rehabilitation dataset and the Urban Land Cover dataset.

D. Sensitivity Analysis

We also study the sensitivity of parameters C, λ, η, and ρ in our algorithm on the Urban Land Cover dataset. Fig. 6 show the results. With the fixed feature dimensions, our method is not sensitive to λ, η and ρ with wide ranges. As for parameter C, when we fix the dimensions, the performance is gradually improved as C increases. When $C > 10^{-3}$, the performance is gradually degraded as C increases. When C is set to 10^{-3}, the performance is the best.

IV. Conclusion

This paper proposed a novel feature transformation method for max-margin based classification. The proposed method aimed to find a low-dimensional feature space to maximize the classification margin of the data and minimize the within class scatter simultaneously. Moreover, we minimized the $l_{2,1}$ norm of the transformation matrix, in order to eliminate the influence of noise and redundant features. Finally, a regularization term was introduced to capture the correlations among multiple classification tasks to help to learn discriminative features. Extensive experiments on publicly available benchmarks demonstrated the effectiveness of the proposed method compared to several related methods.



(a) Vary C and reduced dimensions (b) Vary λ and reduced dimensions

(c) Vary η and reduced dimensions (d) Vary ρ and reduced dimensions

Fig. 6. Sensitivity study of the parameters on the Urban Land Cover dataset.

References

[1] E. Makinen and R. Raisamo, "Evaluation of gender classification methods with automatically detected and aligned faces," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 30, no. 3, pp. 541–547, 2008.

[2] M. E. Mavroforakis, M. Sdralis, and S. Theodoridis, "A geometric nearest point algorithm for the efficient solution of the svm classification task," *IEEE Trans. on Neural Networks*, vol. 18, no. 5, pp. 1545–1549, 2007.

[3] D. Cai, X. He, and J. Han, "Spectral regression: A unified approach for sparse subspace learning," in *IEEE International Conference on Data Mining (ICDM)*. IEEE, 2007, pp. 73–82.

[4] Z. Su, S. Li, and H. Zhang, "Extraction of feature subspace for content-based retrieval using relevance feedback," in *Proceedings of ACM Multimedia (MM)*, 2001, pp. 98–106.

[5] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern classification*. John Wiley & Sons,, 1999.

[6] J. B. Tenenbaum, V. De Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.

[7] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.

[8] M. Belkin and P. Niyogi, "Laplacian eigenmaps and spectral techniques for embedding and clustering." in *neural information processing systems (NIPS)*, vol. 14, 2001, pp. 585–591.

[9] X. He and P. Niyogi, "Locality preserving projections," in *Neural Information Processing Systems (NIPS)*, vol. 16, 2004.

[10] X. He, D. Cai, S. Yan, and H. Zhang, "Neighborhood preserving embedding," in *Proceedings of International Conference on Computer Vision (ICCV)*, 2005, pp. 1208–1213.

[11] H. Wang, S. Chen, Z. HU, and W. Zheng, "Locality-preserved maximum information projection," *IEEE Trans. on Neural Networks.*, vol. 19, no. 4, pp. 571–585, 2008.

[12] F. Wang, B. Zhao, and C. Zhang, "Unsupervised large margin discriminative projection," *IEEE Trans. on Neural Networks*, vol. 22, no. 9, pp. 1446–1456, 2011.

[13] V. Vapnik, *The nature of statistical learning theory*. springer, 2000.

[14] A. Majumdar and R. K. Ward, "Robust classifiers for data reduced via random projections," *IEEE Trans. on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 40, no. 5, pp. 1359–1371, 2010.

[15] S. Paul, C. Boutsidis, M. Magdon-Ismail, and P. Drineas, "Random projections for support vector machines," in *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2013, pp. 498–506.

[16] K. R. Varshney and A. S. Willsky, "Linear dimensionality reduction for margin-based classification: High-dimensional data and sensor networks," *IEEE Trans. on Signal Processing*, vol. 59, no. 6, pp. 2496–2512, 2011.

[17] S. Nikitidis, A. Tefas, and I. Pitas, "Maximum margin projection subspace learning for visual data analysis," vol. 23, no. 10, pp. 4413–4425, 2014.

[18] M. Qian and C. Zhai, "Robust unsupervised feature selection," in *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 2013, pp. 1621–1627.

[19] A. Evgeniou and M. Pontil, "Multi-task feature learning," *neural information processing systems (NIPS)*, vol. 19, p. 41, 2007.

[20] K. Weinberger and L. Saul, "Distance metric learning for large margin nearest neighbor classification," *Journal of Machine Learning Research*, vol. 10, pp. 207–244, 2009.

[21] C. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery*, vol. 2, pp. 121–167, 1998.

[22] A. Tsanas, M. Little, C. Fox, and L. Ramig, "Objective automatic assessment of rehabilitative speech treatment in parkinson's disease," *IEEE Trans. on Neural Systems and Rehabilitation Engineering*, vol. 22, pp. 181–190, 2014.

[23] Y. Zhang and D.-Y. Yeung, "A convex formulation for learning task relationships in multi-task learning," in *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2010, pp. 733–742.

[24] B. Johnson and Z. Xie, "Classifying a high resolution image of an urban area using super-object information," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 83, pp. 40–49, 2013.

[25] D. Cai, X. He, X. Wu, and J. Han, "Non-negative matrix factorization on manifold," in *IEEE International Conference on Data Mining (ICDM)*. IEEE, 2008, pp. 63–72.

[26] M. Qian and C. Zhai, "Robust unsupervised feature selection," in *Proceedings of International Joint Conferences on Artificial Intelligence (IJCAI)*, 2013, pp. 1621–1627.

[27] D. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," *Mathematical Programming*, vol. 45, pp. 503–528, 1989.

[28] Q. Gao, J. Ma, H. Zhang, X. Gao, and Y. Liu, "Stable orthogonal local discriminant embedding for linear dimensionality reduction," *IEEE Trans. on Image Processing*, vol. 22, no. 7, pp. 2521–2531, 2013.

[29] N. Parrish and M. Gupta, "Dimensionality reduction by local discriminative gaussians," in *Proceedings of International Conference on Machine Learning (ICML)*, 2012, pp. 559–566.

[30] J. Chen, J. Ye, and Q. Li, "Integrating global and local structures: a least squares framework for dimensionality reduction," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition(CVPR)*, 2007, pp. 1–8.

[31] D. Cai, X. He, K. Zhou, J. Han, and H. Bao, "Locality sensitive discriminant analysis," in *Proceedings of International Joint Conferences on Artificial Intelligence (IJCAI)*, 2007, pp. 708–713.

[32] Q. Gu, Z. Li, and J. Han, "Joint feature selection and subspace learning," in *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, vol. 22, no. 1, 2011, p. 1294.